Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

### Cover Story
Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

### Featured Articles
Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

### Contact the Editor
To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

### Archives
Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

### Bookmarking
We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

***On behalf of all of us at Intel Developer Update, welcome to the future of the PC platform!***

# Table of Contents

(Click on page number to jump to articles)

# Cover Story

## I/O Building Blocks Keep Pace with Technology

Scott A. Goble
Director of Software Engineering
I/O Products Division
Intel Corporation

### Overview

Scan the technology section of your favorite newsstand and you'll see countless headlines proclaiming the latest step in the march to higher megahertz in processors. The industry roadmap to faster system buses and memory architectures is also well underway. But what about system I/O?

As platform technologies continue to evolve, I/O architectures are likewise beginning to scale to better performance. The PCI bus is moving to PCI-66 and PCI-X, and Ultra3 SCSI is on the way. Switched fabric I/O technologies are also growing in importance, including Fibre Channel and the proposed InfiniBand* architecture.

The growing importance of the Internet in the day-to-day operations of business means that I/O-dependent storage and networking solutions need to scale with these emerging I/O technologies. To help meet this requirement, Intel is announcing a family of I/O building blocks to keep pace with the evolution of I/O technology.

### Intel® I/O Building Blocks

The first members of Intel's I/O building block family are designed to make advanced RAID (redundant array of independent disks) affordable enough to fit a wide spectrum of platforms. Intel® Integrated RAID building blocks enable OEMs, independent hardware vendors, and system integrators to bring the performance and reliability of hardware-based Intelligent RAID to the masses with cost-effective implementations for entry-level servers, workstations, and even performance desktop PCs.

The building blocks include Intel® I/O processors, Intel® Integrated RAID reference designs, and Integrated RAID software that can help OEMs and independent hardware vendors deliver high-performance RAID solutions at affordable price points. In addition, Intel® Integrated RAID controllers provide a quick way to integrate RAID subsystems with selected Intel-based servers.

Let's take a high-level overview of the changes that are coming in I/O technology, and then focus on how Intel Integrated RAID building blocks can help system I/O keep pace.

### PCI Evolves

The PCI local bus is ubiquitous in PC-based platforms because of its processor independence, low pin-count interface, and scalability up to 64-bit I/O performance. While preserving the features that have contributed to its popularity, PCI continues evolving. PCI version 2.2 includes hot-plug capability implemented on the host, and introduces PCI Power Management.

PCI-66 doubles the performance of PCI, to 266 Mbytes/sec, with backward compatibility to PCI. PCI-66 is now beginning to be deployed in volume. PCI-X is an even higher performance extension to the PCI local bus standard, designed to meet the increased I/O demands of Fibre Channel, Gigabit Ethernet, and Ultra3 SCSI. PCI-X is also backward-compatible with PCI devices. PCI-X capabilities include 133 MHz bus speed and 64-bit bandwidth, enabling up to 1 Gbytes/sec throughput. The backward-compatibility of PCI-66 and PCI-X means that the industry is likely to see a mixture of PCI implementations in the near term.

### SCSI Gets Faster

SCSI (Small Computer System Interface) also continues to evolve to keep pace with the growing I/O demands of PCs, workstations, and servers. SCSI is versatile and cost-effective bus used to connect peripherals including CD-ROM drives, scanners, tape drives, and hard disks. Its data path is growing wider and transfer speeds are getting faster to satisfy the requirements of leading-edge systems. The SCSI roadmap points the way to throughput of up to 640 Mbytes/sec with other features for improved manageability.

Ultra3 SCSI is the most advanced SCSI implementation now available. It provides backward compatibility with other SCSI implementations and features a data transfer rate of up to 160 Mbytes/sec. The SCSI roadmap includes future generations: Ultra320 and Ultra640 SCSI, formerly known as Ultra4 and Ultra5 SCSI. Ultra320 SCSI supports data transfer rates of up to 320 Mbytes/sec.

### Fibre Channel

Fibre Channel is a high-speed transmission technology now widely used in storage area networks where multiple hosts connect with storage subsystems. Fibre Channel enables 1 Gbit/sec data transfers and it maps to common transport protocols, including SCSI and IP. One of the big advantages of Fibre Channel is that it merges networking and high-speed I/O in a single highly scalable connectivity technology that enables concurrent communications among workstations, mainframes, servers, attached storage systems, and peripherals.

Unlike Gigabit Ethernet, Fibre Channel is designed to be a transport service that is protocol-independent, providing a single technology for storage, networks, audio, and video. For demanding applications, Class 4 Fibre Channel provides guaranteed delivery and Gigabit bandwidth in addition to fractional bandwidth quality of service. Fibre Channel is a good solution for applications requiring a combination of high I/O performance, reliability, and scalability.

### InfiniBand* Architecture

The InfiniBand architecture is a channel-oriented, switched fabric, serial point-to-point link I/O architecture designed to meet requirements for cost-effective I/O with improvements in reliability, scalability, and performance. InfiniBand provides a single interface for multiple types of I/O and effectively takes I/O expansion outside the server, eliminating slot limits and bandwidth limitations while freeing the CPU to work on applications. InfiniBand architecture supports bandwidth ranging from 500 Mbytes/sec to 6 Gbytes/sec.

High performance is derived from an I/O engine that is directly coupled to host memory, replacing today's shared bus architecture with a fabric of switchable point-to-point links. This approach removes the CPU from the I/O subsystem. The CPU can then communicate with peripherals asynchronously, and the I/O channel engine is responsible for moving data to and from main memory. The bus functions as a switch, enabling point-to-point links to scale with improvements in the performance of CPUs, memory, and peripheral devices.

Due to essentially unlimited bandwidth, the InfiniBand architecture will help align I/O performance with Intel's CPU roadmap, including 32-bit and 64-bit processors. For increased reliability, InfiniBand supports separate fault domains for the CPU complexes and the I/O units while supporting reliable connection mechanisms, data integrity, and fault tolerance. The failure of any unit in the fabric does not impact the remaining nodes.

The draft specification for InfiniBand is currently in preparation, and a final release is tentatively scheduled for early 2000. The first InfiniBand-compliant products should appear during 2001.

### Intel's I/O Roadmap

Intel's I/O processor roadmap scales with the increasing performance of Intel-based server architectures and future I/O bus technologies. Intel is now extending its I/O processor family to support PCI-66. Core processor performance is also on the way up. Future I/O processors, including Intel's next-generation StrongARM* processor architecture implementation, will move core processor performance up to 600 MHz.

Even more important than core performance, however, is the ability of Intel's new I/O processors to scale with the increasing data throughput of emerging I/O technologies, including PCI-X, Ultra3 SCSI, Fibre Channel, and InfiniBand. Intel's I/O processor roadmap will feature deeper caches, improved memory controller interfaces, and integrated bridges, all of which will support the transition to faster I/O implementations. High levels of integration have the additional benefit of reducing chip count and costs.

### The Importance of $I_2O$* Architecture

Intel I/O building blocks are compliant with the I2O architecture specification developed by the I2O Special Interest Group (I2O SIG*). I2O architecture simplifies driver development and helps prepare the way for emerging I/O technology initiatives.

The multitude of current operating systems means that separate drivers must be written, tested, integrated, and supported for every unique combination of OS and device. The I2O architecture defines a "split driver" model that makes drivers portable across multiple operating systems and host platforms, significantly reducing the number of drivers required. OS vendors write a single I2O driver for each class of device and device manufacturers write a single I2O driver for each device, which works for any OS that supports the I2O architecture.

The I2O architecture abstracts I/O operations into a set of messages, which permit the host and the I/O device to communicate. By using the message-passing protocol between independent I/O platforms and the host CPU, I2O software relieves the host of interrupt-intensive I/O tasks and greatly improves I/O performance in high-bandwidth environments.

Here's how Biff Traber, Vice President of Products at Santa Cruz Operations Inc., describes the value of I2O technology: "At SCO, we think the time is right for I2O. Intel has developed an architecture that enables the operating system to talk to an intelligent I/O controller over a separate transport layer. This capability provides a natural lead-in to InfiniBand and other I/O initiatives."

### Summary

The industry is entering a period of rapid transition in I/O architectures that will enable I/O throughput to keep pace with the rapid evolution of Intel® Architecture processors, chipsets, memory implementations, and other platform building blocks. PCI-66, PCI-X, and Ultra3 SCSI implementations are coming in the near-term. Fibre Channel provides a storage area network (SAN) architecture now.

The InfiniBand architecture specification will introduce an I/O architecture effectively free of slot and bandwidth limitations. By offering a roadmap of complementary I/O building blocks, beginning with I/O processors, Integrated RAID reference designs, software, and controllers, Intel is supporting the adoption of new I/O technologies while simplifying the task of integration and improving time-to-market for software and hardware developers.

### More Info

For more information, read the Intel® Integrated RAID Technology white paper and the Intel I/O Building Block Family product brief on the Intel Developer Web site. Or visit the IIO Web site.

### Author Bio

Scott A. Goble is the Director of Software Engineering for the I/O Products Division of Intel's Network Processing Group. Before joining Intel, Scott was a software engineer at Honeywell Corporation, where he worked on the development of Electronic Flight Instruments Systems (EFIS) software for private aircraft. During his seven-year tenure at Intel, Scott's responsibilities have ranged from system validation to market development, prior to his return to software engineering. He holds a B.S. degree in Electrical Engineering.

# Columns

## *Inside Looking In*

### The Magic of a TLA

Tim Mostad
Senior Technical Marketing Manager
Intel Corporation

---

**Overview**

We're still all here. Obviously our technological society hasn't imploded and disintegrated due to the so-called Y2K bug, despite—or maybe because of—the billions of dollars of prevention these three characters provoked.

Surely there are yet to come the $5.25M phone bills (100 years * 365.25 days * 24 hours * 60 minutes * $.10) charged to the unfortunate souls that happened to be on a long distance phone call as the clock rolled over on the last second of the millennium. However, this will generate nothing more than a few laughs and red faces before a corrected billing is quickly generated. No, nothing really significant will ever come from this scare. So how could it happen?

The answer is something we've known for decades in the electronics industry. Moreover, it is something that's become somewhat of a science within Intel; something called the Three-Letter Acronym.

An acronym is a word formed using the initial letters of a name or by combining initial letters or parts of a word. They're useful since they compress strings of big words into simple representations that are easy to remember and turn hard-to-understand sentences into efficient sound bites. It's no wonder why engineers like them. In fact, our ever-efficient engineers have shortened Three-Letter Acronym to TLA. This bit of recursion is somehow all the more satisfying to us nerdy types, as if it serves as proof of symmetry, which we all know is a beautiful mathematical and engineering principle.

From the engineers that gave us CPU, RAM, ROM, PCI, and USB has come yet another catchy, memorable TLA. Fine, so technically Y2K isn't an acronym but most people didn't know that. In fact, many people didn't know that Y2K stood for the problems that could result from inaccurate date calculations based on erroneous hardware and/or software representations of the year portion of the current date or stored past or future dates, roughly speaking. If the media had to say *that* every time the issue came up, do you think the public would have jumped on it with such gusto? I think not.

Actually, I don't know who first coined the term. It was popularized by the media, but its origins were clearly in the engineering world where K stands for 1,024 or as the abbreviation for the metric prefix kilo as in thousands, such as "your software upgrade project is $50K over budget." It was then just a short step from $50K in Q4/99 to Y2K. Fortunately for the people that benefited from our overspending on Y2K fixes, some genius converted the term into a TLA, and the rest is history. Or at least let's hope so.

Even countries that spent next to nothing on Y2K preparedness still manage to exist and the collapse of society apparently isn't imminent, at least from that particular problem. So how did a TLA—and a poorly understood one, at that—manage to generate worldwide anxiety?

I read a commentary the other day that pointed at engineers since “engineers are an easy target.” Engineers are often perceived as magicians practicing some kind of black magic. The creation and perpetuation of a secretly coded language merely serves to further enhance that image. Since we’ve all been taught that “evil” will eventually fail, it is easy to expect the worst from them but it really comes down to the fact that engineers by trade are doers. As I wrote last month, doers eventually fail, since failure is intrinsic to trying. However, Y2K is a shining example of engineers *not* failing.

Sure, there was some well-deserved finger pointing for short sightedness. Corners were cut and compromises made that opened the door to this kind of glitch. Still, aside from the occasional miscreant, engineers don’t choose to design bugs into their products. As the joke goes, “There comes a time in the history of every project when it becomes necessary to shoot the engineers and begin production.” Engineers really want to create perfection, but then the reality of ”get it done” rears its ugly head. No, engineers aren’t the culprits here, just willing participants.

The media, despite the hype, isn’t to blame either. Most of them didn’t understand the problem any better than the people they were talking to. After all, how many critical decisions are based upon the date? Actually very few, and no engineer would knowingly leave a hole in mission-critical code that did, but how could the press know that? Engineers, unable to remember every one of their projects or coding practices, had to remain silent for fear of liability. Granted, problems were found and fixed, and it’s tough to tell exactly how many problems were avoided. How do you measure something that didn’t happen, anyway? No, technically speaking, the *potential* for catastrophic problems existed, and all the reporters had to report it, just as they post reports that speculate on any potential disaster. No, the media were not the culprits here either, just unwitting accomplices.

In the end we only have the magic of a TLA to blame. Y2K became the sum of all unknowable and, at the same time, unforgettable evil. Take note of this and whenever you have anything complex to communicate, spend about 90 percent of your time developing the acronym. As Y2K now proves, you even have some previously unrealized latitude to be creative. At Intel we consider this a WSE (WaterShed Event) in the STC (Science of TLA Creation). Just be sure your TLA has three characters and is something that few can truly understand. And use this magic only for good. The world can only take so much ADO (Acronym Driven Overreaction) about nothing.

### Author Bio

Tim Mostad continues to pursue technical marketing nirvana by applying his 19 years of Intel hardware experience to extending Intel's influence with software and Internet developers. As operations manager in Intel's Developer Relations Division, Tim focuses on the development of broad and efficient enabling processes and infrastructure, primarily through use of the Internet.

## *From the Editor*

Donna Loveland
Managing Editor
Platform Marketing
Intel Corporation

### Overview

Y2K is dead. Long live Y2K!

The much-discussed technology issue has passed and the new year has arrived with the vast majority of us safe and sound and back at work. Our man about the industry, **Inside Looking In** columnist Tim Mostad, shares some wry thoughts on the subject as the rest of us plunge into a technology issue of a different kind—February's *Intel Developer Update*.

Our cover story, "I/O Building Blocks Keep Pace with Technology," takes a look at a major announcement coming out of Intel this month—**I/O Building Blocks**. As developers follow the industry roadmap to faster system buses and memory architectures, what about system I/O? Intel recognizes that the growing importance of the Internet means I/O-dependent storage and networking solutions need to scale with emerging I/O technologies. To help meet this requirement, Intel is announcing a family of I/O building blocks to keep pace. We bring you the details with an industry context.

In this issue you'll find two additional articles on related topics. Fast, affordable design solutions using **Intel® Integrated RAID** is the subject of "Integrated RAID Controller SMU22R Design Kit." For advice on **what to look for in a RAID solution,** see "What's Smart About Intelligent RAID."

As networks grow and expand, what about security, particularly *behind* the firewall? We have answers in "Protect Data from Inside the Firewall," which takes a look at **NIC security adapters**.

If you're developing an EIA-based firmware product, using the **Intel® Applied Computing Firmware System Library** could reduce your cost. The Library provides a clean, simple, 32-bit API for component-level initialization, and it's available free from the Intel developer web site. You'll find the details in "New Firmware Library Reduces Time-to-Market."

If saving time, money, and effort in software development is your goal, check out the new **Intel® Graphics Performance Tool (GPT)**. With GPT you can view low- and high-level performance in real time and analyze the content characteristics of advanced graphics applications without a debug build. We'll tell you how in "Intel® Graphics Performance Toolkit Cuts 3D R&D Time."

For the ultimate in time savings—information to give you a jump on development for the next six months—nothing beats the **Intel Developer Forum**, the computing industry's premier source of tools and training for advanced platform developers. David Barkai, content manager for the upcoming Spring 2000 Conference slated for February 15 – 18 in Palm Springs, summarizes the graduate-level training being offered from Intel senior technical staff and renowned industry experts. As the article recommends, "See What's coming at IDF"—then register to attend.

Between IDF and IDU, you'll get the Intel information you'll need to keep your development A-OK in Y2K. Read on.

### Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

# Departments

## *Applied Computing*

### New Firmware Library Reduces Time-to-Market

Greg Rozzell
Software Technical Marketing Engineer
Intel Corporation

---

### Overview

If the Internet has increased the pace of life in the fast lane for the average businessperson, it has developers of applied computing products moving at the speed of sound. The high rate of technology evolution makes it critical to get to market quickly with innovative, high-performance, connected devices. These sophisticated appliances must also meet increasingly demanding application requirements and tighter budget controls.

Developers regularly face challenges that threaten their ability to meet schedule, feature, and cost goals. One large challenge is in the area of system initialization. Many applied computing platforms don't require the full set of features provided in a standard PC BIOS. The problem is that developing initialization code is difficult and time consuming. The new Intel® Applied Computing System Firmware Library V1.0 (hereafter referred to as Intel® ACSF Library) can help.

### Reduced Time-to-Market and Development Costs

Intel ACSF Library provides a set of libraries and tools that perform core level initialization of the processor, chipset and memory. Intel ACSF Library lets developers implement only the components they need, resulting in faster boot times and a smaller memory footprint. In its minimal configuration, Intel ACSF Library requires only 15 KB of flash (or other nonvolatile) memory space. In addition, Intel ACSF Library can significantly reduce the time and cost associated with developing and validating system initialization modules.

### API for Hardware Abstraction

Currently, Intel ACSF Library provides support for component-level initialization of the Intel® Pentium® II and Celeron™ processors, SDRAM, and the Intel 440BX AGPset. It presents a simple, clean, 32-bit API consistent with common embedded programming practices, regardless which Intel processor or chipset is being utilized. Figure 1 shows a typical initialization stack resulting from an Intel ACSF Library implementation. The API provides an abstraction of the processor and chipset so that if these components change, the developer is no longer required to update any of the "additional initialization code" (as shown) that may have been developed. Instead, the developer downloads only the version of Intel ACSF Library supporting the Intel processor and chipset selected for the new design.

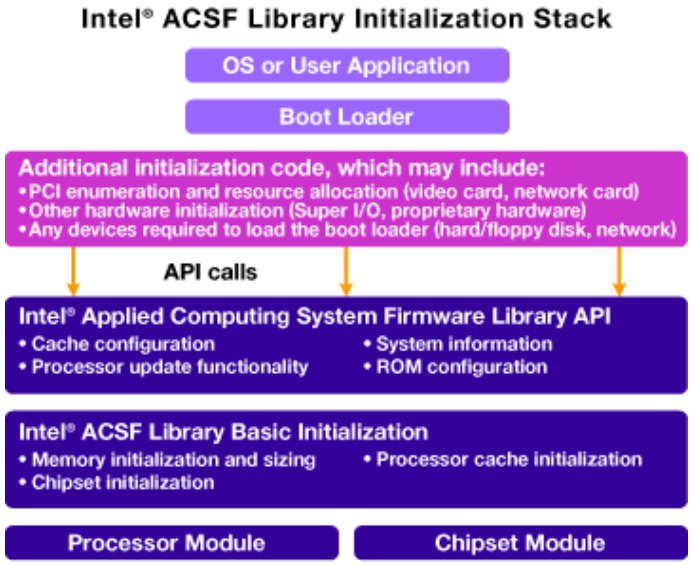


Figure 1

### Available Now, Royalty-free

To get your product to market faster with lower development costs and more time to develop value-added features, download Intel ACSF Library from the Intel software Web site. This tool is provided with a royalty-free run-time license for selected Intel chipset and processor combinations.

### Summary

By using the Intel ACSF Library, developers are able to spend more time on value-added programming and less on creating low-level initialization code. The end result is a product that's easier to differentiate in the highly competitive markets that exist today.

### More Info

For more information about the Intel ACSF Library and Intel Architecture, visit the embedded Intel Architecture Web site. Developers who want the source code can obtain it for a fee by executing a separate license agreement. For more information, contact your local Intel sales representative.

### Author Bio

Greg Rozzell started his career at Orbital Science Corporation in 1989 as an electronics engineer. He joined Intel in -1997 as a software engineer in factory automation, and recently moved to the Software Technology Group. He holds a BSEE and a Master of Computer Science from Arizona State University.

## *Initiatives and Technologies*

### See What's Coming at IDF

David Barkai
IDF Program Manager
Desktop Products Group
Intel Corporation

---

### Overview

If you miss Spring 2000 IDF (Intel Developer Forum) in Palm Springs February 15–17, you could spend the rest of the year catching up. Fall '99 IDF drew more than 3,000 people and triggered over 800 press articles, and the Spring 2000 edition is expected to be even bigger.

More than 20 IDF tracks will provide intensive training from Intel senior technologists and industry experts and more than seven days worth of software technology content.

IDF will begin with board chairman Andy Grove's opening keynote, sharing Intel's vision of the Internet economy. Other keynotes come from Intel business group vice presidents—Albert Yu on the latest advances from his Microprocessor Products Group, Pat Gelsinger of Desktop Products on next-generation platforms, Intel® Architecture Business Group's Paul Otellini on e-Business innovation, and Mark Christensen of Network Communications on the Intel® Internet Exchange™ Architecture and networking.

Here's a closer look at five of the hottest topics coming up at IDF.

### IA-64/Itanium™ Processor

IA-64 has immediate relevance for developers. Being prepared for the Itanium processor launch requires knowledge you can use to implement solutions now, and IDF's IA-64/Itanium™ processor track is the place to get it.

This track focuses on the Itanium processor and Intel's IA-64 roadmap, and more than 80 percent of the course content is devoted to software. Intel technical staff and key industry OEMs, OS vendors, and independent software vendors will share what they've already learned about developing for IA-64. Technical sessions will include preparing software and driver code for IA-64, porting to Windows*, Monterey*, and Linux* operating systems, and developing databases on IA-64. You'll also learn how to use the latest tools and technology for performance optimization and get the latest technical information on resources for IA-64 software development, with a session devoted to the Extensible Firmware Interface (EFI).

### Internet Exchange™ Architecture (IXA)

The Internet is driving rapid changes, and IDF's IXA technical track can help you keep pace. Intel's Network Communications Group V.P. Mark Christensen kicks off the track by scoping out future industry trends. You'll gain a clear understanding of Intel's value proposition for IXA and learn how to use IXA building blocks, including silicon, software, and tools.

Intel customers already using IXA will discuss reference designs and real-world solutions for the Internet service provider, enterprise, and carrier segments. Coverage includes technical details of Intel's network processor, layer 2, 3, and 4 switching, and tools. The track also introduces the next day's IXA Lab. Other networking tracks will cover platform network connectivity and manageability, including broadband and topics in network convergence. The network and communications tracks feature technologies from three of Intel's recent acquisitions—Level One Communications, Dialogic, and IPivot.

### Next-generation IA-32

The evolving IA-32 architecture will continue to deliver connectivity, scalability, speed, accuracy, responsiveness, and enhanced security for users of business and consumer PCs into the 21st century. In this technical track and hands-on lab, you'll get an overview of Intel's next-generation IA-32 microarchitecture and instructions, and key technologies for designing successful power desktop systems utilizing the latest in IA-32 processors and associated chipsets.

Topics to be covered here include performance PC memory architecture, the serial AT Attachment (ATA) storage interface, the PC 2001 Design Guide, power delivery and digital audio, and high-speed design techniques for year-2000 platforms.

A special showcase will highlight the latest concept PCs. The showcase brings together ease of use applications and IA-32 architecture to demonstrate how Intel is bringing interactivity to new areas of the home.

A separate one-day technical track will focus on how to optimize software for the next-generation IA-32 architecture, including methods for porting and optimizing code to the new instructions.

### Intel® e-Business Solutions

The demand for creative and compelling e-Business solutions continues to grow. This track is for anyone who wants to develop e-Business applications with a technical competitive advantage. The focus is on core technology issues, including how Intel Architecture building blocks provide solutions you can use to create a high-availability end-to-end e-Business infrastructure.

This track will also unveil Intel's vision for the third generation of e-Business and provide the tools you need to start developing "3G" e-Business solutions. Sessions will provide detailed coverage of RosettaNet, BroadVision*, one-to-one enterprise relationship management, and end-to-end Linux solutions. It will show how to develop an entire e-Commerce Web site using IA and Web-based problem resolution.

### Software, Software, Software

One of the most important things to remember about IDF is that it's more than a hardware developer's conference. In fact, there'll be so much software-related content at Spring 2000 IDF that a software developer will need to carefully choose which sessions to attend. Fully 80 percent of the three-day IA-64 track is devoted to software. The e-Business track will deliver two days' worth of software-focused content. And one day each is devoted to developing and optimizing software on the next-generation IA-32 and to designing Web sites.

Software and content creation are becoming centrally important in the connected e-Home, where content creates the connected experience. This six-part track discusses emerging video and audio content and content protection, and it presents ways to overcome the barriers to the delivery of modern graphics and video in the home.

### Summary

Beyond the top five topics, the conference will cover mobile and applied computing, with an emphasis on wireless applications. Server, workstation, and systems tracks will provide important tie-ins to IA-64 and next-generation IA-32 technologies. And there'll be a software track on the design of next-generation Web sites.

The Conference's Gold-level sponsors are offering BOFs, special-interest "Birds of a Feather" sessions. The BOFs will include presentations by Analog Devices, Lucent, Novell, Phoenix Technology, Project Monterey, Silicon Image, Texas Instruments, and Viewsonic.

If you'd rather see technology in action, Spring 2000 IDF has demos galore. The Conference has doubled the number of spaces on its exhibition floor, and all slots are filled.

The only other spot at IDF that needs to be filled is yours. Register now, and you'll thank yourself all year long.

### More Info

Visit Intel's IDF Web site for information. You can register for IDF and get details on Palm Springs accommodations as well as further details on keynotes, technical tracks, labs, and demos.

### Author Bio

David Barkai joined Intel in 1996. Before assuming responsibility for IDF Conference content, he worked in the company's Microcomputer Software Lab, where he focused on applications for Intel® Architecture workstations. Prior to that, David specialized in scientific and engineering supercomputing applications and conducted pioneering work on vector processors. He holds a Ph.D. in theoretical physics and has published over 20 papers in the areas of physics, numerical methods, and computer technical applications and architectures.

## *Network and Communications*

### Protect Data from Inside the Firewall

Gregory Youngblood
Product Manager
Network Interface Division
Intel Corporation

---

### Overview

With the rapid growth of e-Business, security is a growing concern for network administrators. While conventional wisdom holds that "security equals a firewall," the truth is that firewalls protect only the front door of the network from external attacks. FBI data shows that the greatest threat to security actually comes from inside the firewall.

The *FBI/CSI 1999 Computer Crime and Security Survey* reports that 31 percent of companies surveyed reported intrusions from outside the firewall. According to the report, *55 percent of these companies experienced unauthorized access by insiders* during the past year, an increase of 10 percent from 1998.

These findings are helping fuel the demand for security inside the firewall. The industry-standard IP Security Protocol (IPSec) developed by the Internet Engineering Task Force offers a strong response. Intel's new PRO/100 S network adapter enables system integrators to offer the performance benefits of hardware-accelerated IPSec security.

### Protect from Inside

IPSec employs data encryption on both the server and the client to protect sensitive data from threats inside the firewall. IPSec protects against "sniffer programs" that intruders can use to steal packets of data, including passwords, off the network. Because IPSec functions at the Network (IP) layer, it's transparent to applications and any network element that operates at Layer 2 or 3.

Due to the demand for inside-the-firewall protection, IPSec is supported as a standard feature of Microsoft's Windows* 2000 operating system. However, packet data encryption algorithms are processor-intensive, and software-based IPSec implementations can burden the CPU and slow application processing on the server and the client. Moreover, the higher packet overhead of host-based configurations puts a drag on network performance.

What's needed is a way to implement IPSec in a way that optimizes server and network performance.

### Hardware-Acceleration Benefits

The coprocessor accelerates IPSec encryption and decryption in hardware, and supports optimum network performance on client and desktop platforms with Microsoft Windows 2000.

Figure 1 shows that Intel's PRO/100 S implementation of the IPSec 3DES protocol results in substantially less CPU utilization and supports higher network throughput, compared to a host-based implementation.

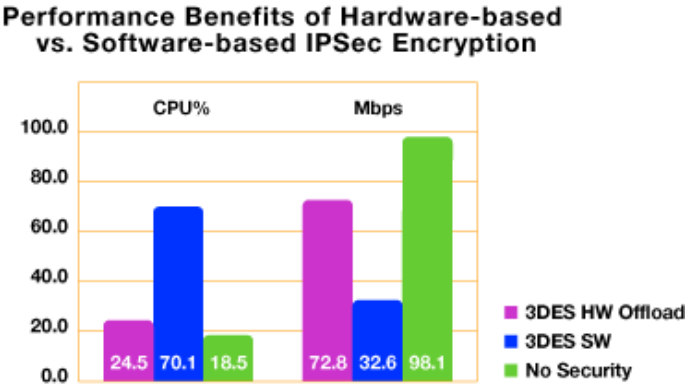


Figure 1

Setup and Configuration Data
Server—Windows* 2000 Advanced Server, Build 2072,
QUAD Proc 500 MHz Pentium® III processor
11 Clients—Windows 2000 Professional, Build 2003,
NetPC Pentium® II processor
Adapter—Intel PRO/100 S
Adapter Driver—Version 4.01.41.0000
Test Benchmark—Chariot version 3.1, Intel hint.scr
Encryption Strength—3DES

### Intel® PRO/100 S Adapter

The Intel PRO/100 S adapter's dedicated intelligent network encryption coprocessor accelerates 56-bit DES (Digital Encryption Standard) and 168-bit 3DES encryption. It also provides MD5 and SHA-1 authentication and data integrity support.

The adapter also includes a wide array of manageability features for connected platforms, including Alert on LAN, Wake on LAN, Wired for Management 2.0, VLANs, and the 802.1p Priority Packet. The adapter includes the Intel® PROSet Utility, Tivoli* agent, and the Intel® LANDesk® Client Manager.

Deployment is fast and easy with Intel® SingleDriver™ technology, a common driver suite that works with all Intel PRO/100 adapters (desktop, server, and mobile) for simplified driver installation and easier maintenance.

### Growing Marketplace Opportunity

Recent statistics from the investment firm ING Baring Furman Selz show the security market segment is growing at an annual rate of 35 percent. At this rate, this market segment will grow from 1997's level of $3.5B to $15.7B in 2002.

Network administrators are concerned with security, and they also demand optimal network and system performance. Implementing IPSec encryption in an intelligent network adapter with hardware acceleration provides a way to meet both objectives.

PRO/100 S cards are inexpensive, easy to deploy, and scalable. Most importantly, many customers are already familiar with Intel "PRO/100 +" network adapters, and this supports a smooth transition to the security-enabled version of this popular card. For only a few dollars more (often less than a 1 percent increase to the cost of a typical PC), customers can get fast, encrypted network connections.

### Summary

With the rise of e-Business, more and more information is being put on corporate LANs and Internet-based VPNs for remote users. Traditional network interface cards (NICs) transmit PC and server data throughout the network in an open fashion, allowing data to be read by anyone inside the firewall who may be "listening" on the wire. Firewalls offer no protection against this growing threat.

Intel's new PRO/100 S security network adapter is an affordable, easy-to-deploy, and scalable solution that developers can use to meet the demand for protection from inside. Software developers should stay tuned to *Intel Developer Update* for further details on interfaces that facilitate the development of third-party security management solutions.

### More Info

Data on the Intel PRO/100 S security-enabled adapter is available on Intel's e-Business Center Web site.

For a summary of statistics from the FBI Computer Crime and Security Survey, see the white paper IP Security: Building Block for the Trusted Virtual Network on Intel's Networking Solutions Web site. Also see Securing the Network in Intel's e-Business Center on the Web.

For more information on IPSec, see IPSec Security and the Corporate Enterprise in the *Intel Developer Update* archive.

Data for developers on the Intel® 82559C Fast Ethernet Controller and the 82594ED Network Encryption Coprocessor is available on the Intel Developer Web site.

### Author Bio

Gregory Youngblood is a product manager in Intel's Network Interface Division. He is responsible for managing all aspects of new product launch activities, including the coordination of marketing, engineering, and manufacturing functions. Gregory holds an MBA and Master's Degree in Systems Engineering from the University of Virginia and a B.S.E.E. degree from Washington University. He is a member of IEEE. Prior to joining Intel, Gregory worked in Lucent Technologies' New Ventures Group as a marketing manager for MPEG-2 compression products.

## *Servers*

### Integrated RAID Controller SMU22R Design Kit

Logan Henriquez
Server RAID Product Marketing Manager
I/O Products Division
Intel Corporation

---

### Overview

As businesses with entry-level and midrange servers adopt Internet-based business models, the "24x7" data availability and reliability of an intelligent RAID subsystem should be a part of every server. As businesses work to improve server reliability and data access, disk storage is often identified as an area of prime vulnerability. The proven solution is Intelligent RAID (redundant array of independent disks), based on controllers that integrate dedicated I/O processors.

The introduction of Intel® Integrated RAID building blocks will help server OEMs and independent hardware vendors (IHVs) overcome the cost barriers which have prevented the wide integration of hardware-based Intelligent RAID in entry-level and midrange Intel-based servers. Intel's Integrated RAID Controller Design Kit (SMU22R) enables vendors to provide affordable controllers that support a spectrum of RAID levels and advanced features for fault tolerance, throughput, and improved manageability.

### The Cost Barrier

Mid-size and smaller businesses engaged in e-Commerce and e-Business confront the same needs for data availability as the largest corporations. In the past, businesses wishing to deploy RAID in their entry-level and midrange servers faced limited options. They could choose the capabilities and performance benefits of a hardware-based RAID solution, with its relatively high price tag. Or they could accept the performance limitations of the more affordable software RAID included with some operating systems. Other businesses, caught between high prices on one hand and performance concerns on the other, have simply gone without RAID. According to statistics released by The Aberdeen Group, fewer than half the Intel-based servers in the sub-$8,000 price category have a RAID subsystem (*Integrated RAID: Building More Storage Functionality into Servers*; Aberdeen Group, Inc. 1999). Intel's Integrated RAID building blocks make intelligent RAID affordable for any size of business and provide a way for OEMs and IHVs to cost-effectively participate in this significant market opportunity.

### What's in the Kit

The Intel® Integrated RAID Design Kit is designed primarily for OEMs and IHVs with manufacturing capabilities in place. The kit provides instructions and schematics needed to develop an Ultra2 LVD (low voltage differential) SCSI dual-channel RAID controller based on the highly integrated data flow architecture of the Intel® i960® RM I/O processor. Kit contents include the SMU22R Reference Design (board and Integrated RAID software), the KMU21 Integrated RAID Controller card, SMU22R schematics, and technical documentation.

The reference design supports RAID Levels 0, 1, 5, and 10, enabling developers to support their customers with RAID implementations featuring the best balance of fault tolerance and data throughput. In addition to the Intel® I/O processor, the design includes an LSI Logic* SYM53C896 SCSI controller, supporting a burst data transfer rate of up to 80 Mbytes/sec. The reference design supports from 32 to 128 Mbytes of ECC SDRAM and features battery back-up.

### Integrated RAID Software

Intel® Integrated RAID software enables developers to offer advanced RAID capabilities that were once available only in high-end controllers. These capabilities include online capacity expansion, online RAID level migration, and adjustable rebuild priority, with no need to power-down or reboot the system. For uninterrupted access to data, Integrated RAID supports hot-plug drives, hot spares, and automatic rebuild capability. To optimize throughput, full and sequential write-back caching is user-selectable for each RAID volume.

HTML-based software simplifies “out of the box” set-up, with instant availability and background initialization. The common look and feel of the browser-based configuration utility helps eliminate user retraining in cross-platform environments and allows the system administrator to manage the system remotely. For OEMs, the Design Kit includes a CD containing manufacturing utilities, diagnostics, schematics, and customizable templates for end-user documentation.

### Summary

Intel’s Integrated RAID Controller Design Kit is one of a family of affordable Integrated RAID building blocks for developers. It’s a fast time-to-market design solution designed primarily for OEMs and IHVs with manufacturing capabilities currently in place. The design supports high levels of RAID performance with a highly integrated Intel I/O processor, an Ultra2 SCSI controller, and up to 128 Mbytes of ECC SDRAM. The kit also includes Intel Integrated RAID software, providing advanced RAID management capabilities that were once reserved for high-end RAID controllers.

Intel’s Integrated RAID building blocks provide multiple levels of integration that enable developers to choose the building blocks that best match their available resources. The Intel Integrated RAID Controller Design Kit provides an easy way for OEMs and IHVs to meet the growing customer demand for reliable, easy to manage RAID solutions that offer ultimate I/O performance for the new connected world.

### More Info

For more details, check out the Intel Integrated RAID Demo and read the article introducing Intel Integrated RAID on the Intel Developer Web site. Or visit the IIO Web site.

### Author Bio

Logan Henriquez is server RAID product marketing manager for Intel's I/O Products Division. Logan's experience includes several years leading the Compaq server storage division product marketing team, which launched first Fibre Channel product, created its RAID product line, and grew the division into a leading storage supplier. Logan also held positions in software development and planning at Intuit and American Express International. He holds a BSIE from Stanford University and an MBA from the University of Chicago.

## What's Smart About Intelligent RAID?

Scott Wells
Product Marketing Manager
I/O Products Division
Intel Corporation

### Overview

Reliability has become the single biggest buzzword in the world of Internet computing. It's easy to see why. With everything from e-Commerce and important business applications to basic file and print services running on server-based networks, reliable access to data is becoming a critical issue for businesses of all sizes.

As the Internet and e-Business applications grow in importance, small and mid-sized businesses are beginning to recognize the value of RAID technology as a way to protect their vital data against the cost of unplanned interruptions caused by disk drive failure. Of the various RAID technologies now available, Intelligent RAID is often the best choice for a server-based business.

Intel's new family of Integrated RAID solutions is helping to make advanced Intelligent RAID affordable enough for deployment in entry-level and midrange servers. As this trend continues, Intelligent RAID subsystems will become more and more pervasive in servers, workstations, and even performance desktop PCs.

### RAID Basics

RAID stands for "redundant array of independent disks" that appear to the system as a single logical drive, and it is becoming an increasingly important feature of Intel-based servers. When you begin to look at a RAID solution, the first thing you realize is that a number of different implementations, or "levels," of RAID are currently available. Each level is designed to meet a specific set of requirements. Once you understand what makes the RAID levels different, you can choose the level that meets the needs of your application.

Before you get started, you also need to know that RAID can be based on three different technologies. Making the smart choice now is the best way to ensure a high level of server reliability right away, with a built-in path to future expansion that can grow along with your business requirements. We'll scope out the levels, then take a closer look at the technologies.

### Getting in on the Right Level

RAID is available in a number of "levels." It helps to understand what the various levels mean. Even more important is knowing the degree of data protection each level can provide, and then matching those features to specific usage requirements of the particular server. Here are some examples:

RAID Level 1, or "disk mirroring," provides a high level of data reliability by essentially duplicating the contents of one disk onto another. While mirroring offers the highest reliability, it also doubles the cost of the storage system.

Other RAID levels store "stripes" of data across multiple drives, which can increase performance. Some RAID systems combine striping with "parity," which improves reliability by calculating data on two drives and storing the result on a third. Parity allows the data to be reconstructed in the event of drive failure. Parity calculations can be processor-intensive.

RAID Levels 3 and 5 are popular because they combine striping with parity to balance reliability, performance, and price. In RAID Level 3, parity results are stored on one dedicated drive, while Level 5 "rotates" parity data across several drives in the array.

RAID Level 3 is a good choice for the "single-thread reads" and small data blocks handled by servers that support multiple users on a network; RAID Level 5 adds more reliability.

Once you have a basic understanding of RAID levels, you need take a closer look at the underlying RAID technology.

### All RAID Is Not Created Equal

There are three basic RAID technologies available today:

- **Host-based RAID** uses a scheme in which the server's CPU must handle all RAID interrupts and RAID processing tasks. This additional overhead can impact I/O throughput and also affect the server's CPU and system-level performance.

- **Hardware-assisted RAID** is a hybrid approach that offloads some, but not all, RAID-related processing tasks. The CPU and other system resources are still required to handle some RAID processing, which can impact system-level performance.

- **Intelligent RAID** features a RAID controller with a dedicated I/O processor. The I/O processor runs the RAID software algorithm that manages the flow of data between the disk subsystem and the server's CPU. Intel® I/O processors feature a highly integrated data flow architecture specifically designed to improve I/O throughput. Intelligent RAID has another important advantage. It is the only technology that off-loads all RAID-related interrupts and processing tasks from the CPU. This frees the CPU for application processing while optimizing system-level performance. For these reasons, Intelligent RAID provides the best combination of reliability and data availability, without compromising throughput or performance.

### Intelligent RAID Is the Smart Choice

If you've compared Intelligent RAID with hardware-assisted (hybrid) RAID and host-based RAID, in the section above, you know that Intelligent RAID supports optimal system-level performance. But its advantages don't end there.

Intelligent RAID also offers more complete system-level protection. Unlike host-based RAID systems, Intelligent RAID makes it easier to protect the server's boot drive, where the operating system resides.

The ultimate reasons to choose Intelligent RAID involve its growing availability and relatively low cost. Some major OEMs are now implementing Intel-based Intelligent RAID controllers directly on the server board, a feature known as ROMB (RAID-on-Motherboard). In addition, a number of leading independent hardware vendors, including AMI, DPT, Mylex, and ICP vortex, now offer Intelligent RAID adapter cards based on Intel I/O processors. And Intel® Integrated RAID building blocks can make it relatively simple to bring an Intelligent RAID adapter to market.

### Summary

As the Internet becomes more pervasive, the mission-critical data storage needs of businesses of all sizes will continue to grow. When you're investing in a RAID system, it's a good idea to plan ahead by selecting a RAID technology that can scale with these growing business requirements. This is an area where Intelligent RAID shines. Dedicated RAID controllers make it easy to add and configure additional disks, and more controllers can be added, as needed, all without impacting server's system-level performance. RAID controllers that feature hot-swap support enable operators to add and configure new drives while the system stays up and running.

By supporting server reliability, availability, and performance, Intelligent RAID can help businesses of all sizes maximize user productivity, minimize downtime, and reduce data recovery costs. It can add up to reduced Total Cost of Ownership for businesses, and opportunities for system designers.

### More Info

For more information about Intel Intelligent RAID building blocks, see the article "Intel Intelligent RAID" on the Intel Developer Web site and check out the interactive Intel Integrated RAID Demo. You can also see the RAID Product Selector and the IIO Web site.

Also be sure to see the articles “Integrated RAID Controller SMU22R Design Kit” and “I/O Building Blocks Keep Pace with Technology” in this issue of *Intel Developer Update*.

### Author Bio

Scott Wells is a product marketing manager within Intel’s I/O Products Division. He currently manages the Intel Integrated RAID Controller KMU21 and the Intel Integrated RAID Design Kit SMU22R, and is responsible for coordinating project planning, marketing, and product launch activities. He holds a B.S. degree in Systems Analysis from Miami University and an MBA from the University of Dayton.

## *Software*

### Intel® Graphics Performance Toolkit Cuts 3D R&D Time

Hersh Reddy
Technical Marketing Engineer
Intel Architecture Content Group
Intel Corporation

---

### Overview

One of the biggest challenges facing 3D graphics developers today is the diversity of hardware available for 3D applications. Developers must debug and run performance tests on a variety of platform configurations, which means incorporating performance analysis and tracking capabilities into graphics applications, along with the ability to output and record statistics. Adding this instrumentation to an application is both tedious and time-consuming, which is why Intel developed the Graphics Performance Toolkit (GPT) v3.0 to help developers benchmark and analyze graphics applications.

### The Intel® GPT

The Intel Graphics Performance Toolkit (GPT) v3.0 is an integrated suite of tools that combines extensive API-level function logging with run-time statistics and a real-time analysis mode. Because GPT is easy to use, most members of the development team, from graphic artists to quality-assurance technicians, can run performance analyses on 3D graphics applications, thereby freeing developers to spend their time fine-tuning and optimizing software. GPT modifies (or instruments) the API DLLs, Ddraw*.dll for DirectX* 7.0, and Opengl*32.dll for OpenGL* 1.1 to log and analyze the API calling behavior and performance of an application. Instrumenting DLLs keeps source code from getting cluttered with countless repetitions of *#ifdef DEBUG*. Since GPT modifies the API DLLs directly, developers can actually analyze release-candidate executables without source code.

Users define the level of detail to be included in the analysis log. GPT can act simply as an instantaneous frame rate counter or it can log every API level event, even listing each vertex passed to each *DrawPrimitive* function call. To save disk space, the log time can be configured in a number of ways: the application activity can be logged over its entire lifetime, by start and stop keys, or even by specifying a time interval in which to activate logging. The log information can then be configured to show a number of statistics from render state changes per frame to triangles submitted to the graphics API. Information can be viewed graphically or with a spreadsheet.

GPT provides four analysis tools that are accessed through a common user interface: the Frame Rate tool, the Performance Analysis tool, the Advanced Analysis tool, and the Real-time Analysis tool.

### Frame Rate Tool

The Frame Rate tool logs instantaneous frame rate information. After a run, the user can choose to display a graph in GPT that shows the frame rate according to the frame number starting from the first logged frame. (See Figure 1.) GPT can also display an instantaneous frame rate graph at run time that runs unobtrusively over the application. The tool offers several different ways of defining just exactly what a “frame” is. For example, in DirectX, frames can be defined as the code run between two consecutive calls to *BeginScene* or *Flip*. Or, alternatively, frames can be delimited by consecutive *Blts* to the primary surface. The frame delimiter (used by the graphics API) is chosen in the GPT user interface**.** The DirectX API frame delimiters are *BeginScene*, *Flip*, and *Blt*. The OpenGL version of GPT similarly uses *wglSwapBuffers*, *glDrawPixels*, *glClear* (*GL_DEPTH_BUFFER_BIT*), *glClear* (*GL_COLOR_BUFFER_BIT*), or *glClear* (*GL_DEPTH_BUFFER_BIT* | *GL_COLOR_BUFFER_BIT*).

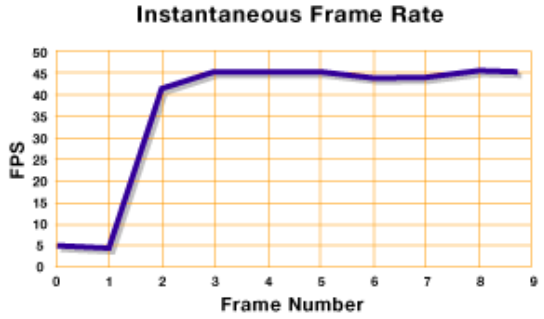


**Figure 1: Graphed Frame Rate Information**

### Performance Analysis Tool

This tool logs the time spent in each API function as well as in the total calls made. Performance analysis information can be displayed on a frame-by-frame basis or as a sum total for the entire logging session. The Performance Analysis tool also lets the user disable lighting and/or texturing to determine the effect these features have on graphics performance. Setting GPT to serialize all graphics calls tests CPU and 3D-accelerator concurrency. Serializing graphics calls should slow the application down significantly, indicating that concurrency is at an optimal level. If the application is not slowed down, further optimization is required to distribute the load more evenly between the CPU and the graphics card. Running GPT with NULL API activated effectively deletes all graphics API function calls to determine whether the graphics code is the performance bottleneck, as opposed to non-graphics routines.

### Advanced Analysis Tool

The Advanced Analysis tool can log a prodigious amount of run-time data. (For example, Figure 2 shows GPT's graph of triangle list length distribution per frame.) This includes parameters to each API function call as well as return values. As mentioned earlier, GPT can even be configured to record every vertex passed to the graphics API. This level of detail is generally not needed but could be useful in identifying hard-to-find bugs in lighting and transformation code. A more generically useful feature is the extensive texture logging, which details what the textures are doing in every frame. Texture logging includes details such as the number of texels referenced, the number of texture state changes, or the number of textures loaded into the display device. The Advanced Analysis tool can also graph the total primitives submitted to the API and show exactly when the hardware is being stressed the most.

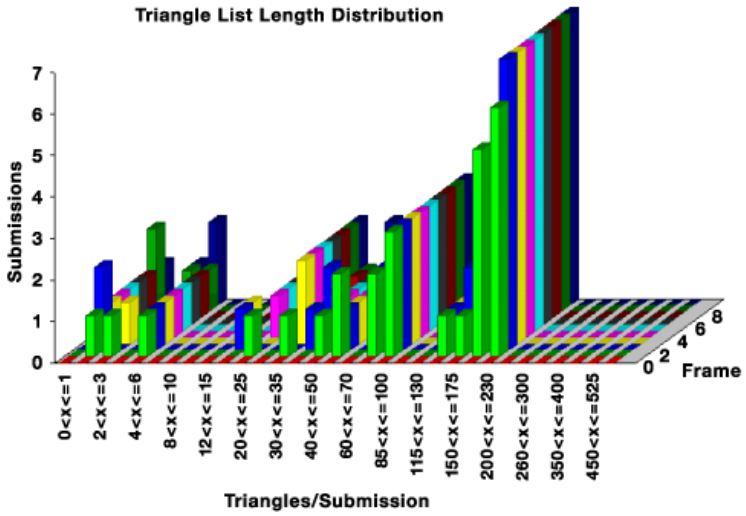


**Figure 2: Triangle List Length Distribution Per Frame**

### Real-time Analysis Tool (Quad View)

This tool displays four different views of the same graphic scene simultaneously: wireframe, non-textured, overdraw, and normal. Each view can be toggled to display in full-screen mode for closer analysis. The wireframe view in the bottom left section of Figure 3 shows the underlying frame of the graphic. The non-textured view in the bottom right section shows the graphic flat and shaded without texture. The overdraw view on the top right shows the transparent graphic with overlay and shading. Normal view, top left, is the fully textured 3D graphic, or what the end user would see.

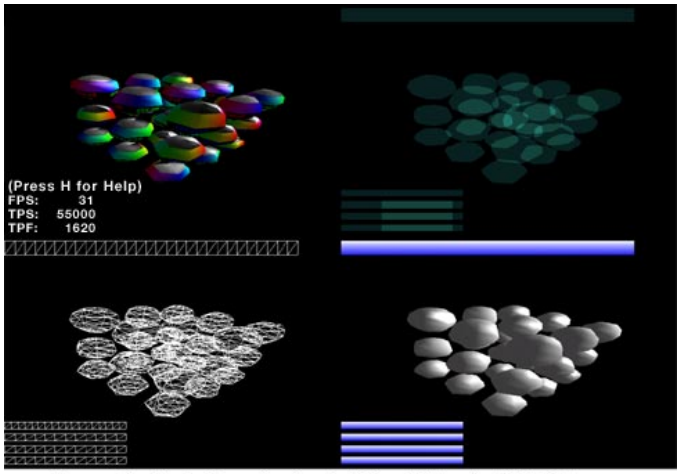


**Figure 3: Real-time Analysis (Quad View) of 3D Spheres**

The Real-time Analysis tool lets developers and non-technical team members analyze 3D graphics applications using these four different views. For example, a level tester who finds the application slowing down at a certain point can toggle to overdraw view in full-screen mode and check to see whether the scene complexity is too high or whether the portal culling has an error. Bugs in an object importer are also obvious when viewing 3D information in wireframe.

### Summary

The Intel Graphics Performance Toolkit can be an enormous time-saver at crucial points in the development cycle, especially during the graphics display system design process. By using GPT, developers can leave some performance analysis work to others on the development team. This frees developers to optimize and fine-tune software for the variety of hardware platforms on the market today.

### More Info

GPT v3.0 for OpenGL and DirectX is now available for Windows* 98. For more information or to download a 14-day evaluation copy, see the Intel Graphics Performance Toolkit Web site.

### Author Bio

Hersh Reddy joined Intel in October 1999. He is currently a technical marketing engineer in IACG and is involved with gathering, researching, and disseminating scalable solutions for real-time 3D Graphics, AI, and Physics. Previous to his work at Intel, he worked as a programmer in the video game industry, shipping one title (Revenant for Cinematix Studios and Eidos Publishing). He holds a B.S. degree in Computer Science from Cornell University in New York.

—End of Intel Developer Update Magazine Issue 5—